

accessing a single one of said physical disk units which has a minimum number of operations from among the disk units storing the designated logical volume and outputting a request to only said accessed single minimum waiting physical disk unit;

incrementing the number of operations of said accessed physical disk unit in accordance with a request on said operation ; and

decrementing the number of operations of a physical disk unit whose operation has been completed, in accordance with an end of said operation, wherein each of said physical disk units performs requested operations in a queued order, and

wherein said program further comprises referring to a memory storing a table indicating a correspondence of the plurality of physical disk units and said designated logical volume to select single physical disk unit on which said designated logical volume is allocated in accordance with said designation of said designated logical volume by a high rank apparatus.

63
Cancel

Add
H 2

REMARKS

In the final Office Action mailed March 8, 2002, claims 1, 3-5, 7, 9-11, and 13-16 were rejected under 35 USC §102(e) as being anticipated by Tanaka et al. (U.S. Patent No. 5,542,064).

Claims 4 and 10 are canceled. Claims 1, 7, and 13-16 are amended. A Version with Markings to Show Changes Made to the claims is enclosed herewith.

Claims 1, 3, 5, 7, 9, 11, and 13-16 are pending, of which claims 1, 7, and 13-16 are independent. Claims 3 and 5 depend, either directly or indirectly from claim 1 and claims 9 and 11 depend, either directly or indirectly from claim 7.

The rejections of remaining, pending claims 1, 3, 5, 9, 11, and 13-16 are respectfully traversed.

The Tanaka apparatus provides multiple writing of identical data in a secondary storage device (see column 2, lines 25-52). That is, Tanaka discloses a RAID system in which a read/write command is given to all the disk drives storing identical data (see, column 2, lines 1-13). Therefore, Tanaka discloses that it is necessary to select multiple disk drives and output a

request to all of the selected multiple disks (see column 8, lines 1-47). Further, in order to prevent waiting, Tanaka discloses canceling commands to selected disk drives except for the first connected disk drive (see column 8, lines 48-55).

In contrast to the Tanaka apparatus, the present invention balances waiting requests between disk drives, by selecting the single disk drive with the lowest wait and outputting a request to only selected single disk drive.

Independent claims 1, 7, 13, 14, 15, and 16 each recite the foregoing features of the present invention. More particularly, each of independent claims 1, 7, 14, and 15 (using the recitation of claim 1 as an example) of the present application recites "selects the single physical disk unit from among the disk units storing the designated logical volume which has a minimum number of operations based on the comparison" and "outputs a request to only the selected single minimum waiting physical disk unit".

In addition, each of independent claims 13 and 16 of the present application recites (using the recitation of claim 13 as an example) "accessing one of the first and the second of the redundant logical volumes based on a minimum number of the numbers of operation respectively requested of each of the physical disk units storing the redundant logical volumes based on the counting, and outputting a request to only the accessed single minimum waiting physical disk unit".

Tanaka does not disclose or suggest the foregoing features of the present invention.

Moreover, the above-mentioned dependent claims recite patentably distinguishing features of their own. For example, claim 3 (depending from claim 1) recites "'a resource manager circuit determining one of the plurality of physical disk units to be accessed in accordance with said number of operations in said memory in response to a transfer request from said channel adapter circuit, and requesting said device adapter circuit to perform an operation accessing said determined physical disk unit".

Withdrawal of the foregoing rejections of claims 1, 3, 5, 7, 9, 11, and 13-16 as being anticipated by Tanaka et al. is respectfully requested.


If there are any formal matters remaining after this response, the Examiner is requested to telephone the undersigned to attend to these matters.

If there are any additional fees associated with filing of this Amendment, please charge the same to our Deposit Account No. 19-3935.

Respectfully submitted,

STAAS & HALSEY LLP

Date: June 10, 2002

By: 
Gene M. Garner II
Registration No. 34,172

700 Eleventh Street, NW, Suite 500
Washington, D.C. 20001
(202) 434-1500

VERSION WITH MARKINGS TO SHOW CHANGES MADE

IN THE CLAIMS:

Please CANCEL claims ~~9~~ and 10. ✓ ✓

Please AMEND the following claims:

1. (SEVEN TIMES AMENDED) A RAID apparatus comprising:

a plurality of physical disk units storing a plurality of copies of each of logical volumes;
and

a disk controller accessing any of the physical disk units which stores a designated logical volume to thereby access said designated logical volume,

said disk controller including:

a memory storing the number of operations requested to each physical disk unit,
for each physical disk unit, and

control means for accessing one of said plurality of physical disk units which stores the designated logical volume, in accordance with said number of operations,

wherein said control means compares numbers of operations corresponding to a plurality of physical disk units which store said designated logical volume with each other, selects the single physical disk unit from among the disk units storing the designated logical volume which has a minimum number of operations based on the comparison, and outputs a request to only the selected single minimum waiting physical disk unit,

wherein said control means increments the number of operations of said selected physical disk unit in accordance with a request for said operation and decrements the number of operations of a physical disk unit whose operation has been completed, in accordance with an end of said operation,

wherein each of said physical disk units performs requested operations in a queued order, and

wherein said memory stores a table indicating a correspondence of the plurality of physical disk units and said designated logical volume; and said control means refers to said memory with said designated logical volume to select said single physical disk unit on which

said designated logical volume is allocated in accordance with said designation of said designated logical volume by a high-rank apparatus.

3. (AS TWICE AMENDED) The RAID apparatus according to claim 1, wherein said control means includes:

- a channel adapter circuit performing interface control with said high-rank apparatus;
- a device adapter circuit accessing said physical disk units in accordance with a requested operation; and
- a resource manager circuit determining one of the plurality of physical disk units to be accessed in accordance with said number of operations in said memory in response to a transfer request from said channel adapter circuit, and requesting said device adapter circuit to perform an operation accessing said determined physical disk unit.

5. (AS ONCE AMENDED) The RAID apparatus according to claim 3, wherein said memory stores status information indicating statuses of said physical disk units; and

said resource manager circuit refers to said status information to determine whether those physical disk units which store said designated logical volume are normal and selects a normal physical disk unit.

7. (SIX TIMES AMENDED) An access control method for a RAID apparatus comprising a plurality of physical disk units storing a plurality of copies of each of logical volumes, and a disk controller accessing any physical disk unit which stores a designated logical volume to thereby access said designated logical volume, said method comprising:

determining a plurality of physical disk units which store a designated logical volume; and

selecting from among the determined disk units storing the designated logical volume one of said determined physical disk units in accordance with the number of operations requested to said physical disk units, said selecting comprising:

comparing said numbers of operations of a plurality of physical disk units which store

said designated logical volumes with each other,

accessing the single physical disk unit which has a minimum number of operations based on the comparison and outputting a request to only said selected single minimum waiting physical disk unit,

incrementing the number of operations of said accessed physical disk unit in accordance with a request for said operation, and

decrementing the number of operations of a physical disk unit whose operation has been completed, in accordance with an end of said operation,

wherein each of said plurality of physical disk units performs requested operations in a queued order, and

wherein said selecting further comprises referring to a memory storing a table indicating a correspondence of the plurality of physical disk units and said designated logical volume to select a physical disk unit on which said designated logical volume is allocated in accordance with said designation of said designated logical volume by a high-rank apparatus.

9. (AS THRICE AMENDED) The access control method according to claim 7, wherein said determining step determines said plurality of physical disk units in response to a transfer request from said high-rank apparatus; and

said selecting includes:

requesting an operation for accessing said physical disk unit determined in accordance with said number of operations, and

accessing said physical disk unit in accordance with said requested operation.

11. (AS TWICE AMENDED) The access control method according to claim 7, wherein said selecting includes:

referring to status information to determine indicative of statuses of said physical disk units, stored in said memory, to determine whether those physical disk units which form said designated logical volume are normal; and

selecting a normal physical disk unit.

13. (SIX TIMES AMENDED) A RAID apparatus comprising:

physical disk units storing redundant logical volumes, a first of the redundant logical volumes being stored on one of the physical disk units, and a second of the redundant logical volumes being stored on another of the physical disk units; and

a disk controller counting numbers of operations respectively requested of each of the physical disk units and accessing one of the first and the second of the redundant logical volumes based on a minimum number of the numbers of operation respectively requested of each of the physical disk units storing the redundant logical volumes based on the counting, and [outputs] outputting a request to only the accessed single minimum waiting physical disk unit,

wherein said disk controller increments the number of operations of an accessed physical disk unit in accordance with a request for said operation and decrements the number of operations of an accessed physical disk unit whose operation has been completed, in accordance with an end of said operation,

wherein each of said physical disk units performs requested operations in a queued order, and

wherein said disk controller refers to a table indicating a correspondence of the plurality of physical disk units and said redundant logical volumes to select the accessed single physical disk unit on which one of said redundant logical volumes is allocated in accordance with said designation of said one of said redundant logical volumes by a high-rank apparatus.

14. (FIVE TIMES AMENDED) A RAID controller accessing one of a plurality of physical disk units storing a plurality of copies of each of logical volumes to thereby access a designated logical volume, comprising:

a memory storing a number of operations requested [to] of each physical disk unit corresponding to each physical disk unit; and

a controller comparing said numbers of operations corresponding to a plurality of

physical disk units which store [said] a designated logical volume with each other, and selecting single one of said plurality of physical disk units which has a minimum number of operations from among the plurality of physical disk units storing the designated logical volume based on the comparison and outputting a request to only said selected single minimum waiting physical disk unit,

wherein said controller increments the number of operations of said selected physical disk unit in accordance with a request for said operation and decrements the number of operations of a physical disk unit whose operation has been completed, in accordance with an end of said operation, wherein each of said plurality of physical disk units performs requested operations in a queued order, and

wherein said memory stores a table indicating a correspondence of the plurality of physical disk units and said designated logical volume; and said controller refers to said memory with said designated logical volume to select a physical disk unit on which said designated logical volume is allocated in accordance with said designation of said designated logical volume by a high-rank apparatus.

15. (FIVE TIMES AMENDED) A balancing access method for a RAID apparatus comprising a plurality of physical disk units storing a plurality of copies of each of logical volumes, comprising:

comparing numbers of operations of a plurality of physical disk units which store a designated logical volume with each other;

selecting a single one of said physical disk units which has a minimum number of operations from the disk units storing the designated logical volume based on the comparison and outputting a request to only said selected single [mimimum] minimum waiting physical disk unit;

incrementing the number of operations of said accessed physical disk unit in accordance with a request on said operation; and

decrementing the number of operations of a physical disk unit whose operation has been completed, in accordance with an end of said operation, wherein each of said physical disk units performs requested operations in a queued order, and

wherein said selecting further comprises referring to a memory storing a table indicating a correspondence of the plurality of physical disk units and said designated logical volume to select said single physical disk unit on which said designated logical volume is allocated in accordance with said designation of said designated logical volume by a high-rank apparatus, and wherein said balancing access method auto-adjusts loads between the physical disk units.

16. (FIVE TIMES AMENDED) A storage medium of a RAID apparatus storing a program, said program which when executed by a computer causes the computer to execute processes comprising:

comparing numbers of operations of a plurality of physical disk units which store a designated logical volume with each other;

accessing a single one of said physical disk units which has a minimum number of operations from among the disk units storing the designated logical volume and outputting a request to only said accessed single minimum waiting physical disk unit;

incrementing the number of operations of said accessed physical disk unit in accordance with a request on said operation ; and

decrementing the number of operations of a physical disk unit whose operation has been completed, in accordance with an end of said operation, wherein each of said physical disk units performs requested operations in a queued order, and

wherein said program further comprises referring to a memory storing a table indicating a correspondence of the plurality of physical disk units and said designated logical volume to select single physical disk unit on which said designated logical volume is allocated in accordance with said designation of said designated logical volume by a high-rank apparatus.